Hearing Smarter: How AI is Transforming Hearing Aids

Presenter: Joshua M. Alexander, PhD, CCC-A

It's my pleasure to welcome back Doctor Joshua Alexander. Doctor Alexander comes to us from Purdue University. If you'd like to learn more about him and his background, you can read his bio on the course registration page, but I'll go ahead and hand the mic over to you. Doctor Alexander, very good. Thank you for the opportunity to give you a presentation on AI and hearing aids.

So a lot of what you've, obviously, AI is everywhere, which kind of prompts this discussion here, just kind of, and depending upon how much you've been exposed to, some of this material may be a review, some of it may be new to you. Hopefully, a lot of what you've been exposed to probably has been coming straight from the manufacturers because they have some of the inside knowledge of what's they're doing to their products. And what I'm going to present to you is kind of going to be a summary overview of what's available across the products. So I'll be talking about a lot of different technology today. Not all of this technology exists in a single hearing aid.

And throughout this presentation, I will refer to specific examples that manufacturers are using, but I will not be referring to those manufacturers by name. So with that, let's go ahead and get started.

Here's my disclosures. Here are the learning outcomes that we will try to achieve today with today's talk that we need to take a few steps and just think about, obviously, AI, again, is everywhere in our world. We're being exposed to it for multiple fronts. Just like when digital hearing aids first came out, we have to think about is there actually even a need for this? What is the need for AI in hearing aids?

Or is this just kind of another attractive feature that maybe doesn't add a whole lot of value added benefit? And so, kind of referring back to a statement that I made a couple years ago, if we think about what hearing aids do, we know hearing aids make sound loud. Making sound loud is easy. And for those of us who try to communicate, to lay individuals who don't have an understanding of why hearing aids are so complex, why they don't solve all of our problems, why hearing aids cost a lot of money, they conceptualize hearing aids as just being pure amplifiers. Well, we can make sound loud.

That's the easy part. And one of the most important jobs, the hearing aid, is to control the amount of amplification, to control the amplification process by attenuating, not turning up, but by attenuating the right sounds at the right frequencies, the right times, the right direction of arrival and you know, the easiest example of this doesn't require a lot of intelligence is by dynamic range compression, right? So for low inputs, we have apply a lot of gain. As the input goes up, input level goes up, we turn the gain down. And the operative word here is control.

This control implies some decision making that has to be done somewhere and based on some amount of information. And so the easiest way to do this is to create rules, create rules based on the overall level, the signal to noise ratio at different frequencies, the modulation features to tell us how much is speech, how much is noise, that kind of thing. And the more variables that we can consistider in this decision making process, presumably the better the resultant output will be for the patient. However, as humans, we are ultimately limited by our ability to integrate a large number of variables. We can talk about a discrete number of variables, but once we start expanding the possibilities of things to consistider in our decision making process, that's where the limitations fall in, especially if these things are changing rapidly over time and to be able to rapidly adjust.

And so, even if we get the control process right, on average, for the average person, there are still challenges that are associated with individual users and with individual situations. So we have to make, when designing these, at the onset, we have to make assumptions about what's meaningful to the user and to what the average, average person's ear and brain need to process the information, and then what is meaningful to the user in a particular situation. It's not always speech, and it doesn't always sound, doesn't always arrive the thing that they want to hear from the front.

Now, this is going to be, again, this big, broad overview, but then we'll take the deep dive. So what is AI? Artificial intelligence refers to the creation of systems that can perform tasks typically requiring human intelligence, such as learning, decision making, and pattern recognition. So you probably could have told me that before we started this thing. Is the AI systems the key there is they improve over time, they improve their performance, and they do this by analyzing large amounts of data, identifying patterns in that data, and then making predictions based on that data.

There are two different, there are multiple subtypes of AI, but there's two that you see a lot in the hearing aid literature. There's machine learning. This is a subset of AI, where the systems learn explicitly, learn from the data, improve over time without being explicitly programmed. We use algorithms to find patterns in the data. Once the system learns from the data, they can make decisions or predictions.

Then we have deep learning. This is one that you've probably seen as well, deep learning, deep learning, neural networks. We'll talk a little bit about that. Deep learning is a subtype of machine learning. It uses neural networks.

And I'll explain what this means in a moment with multiple layers, hence the word deep, to process complex information. They're named neural networks because they mimic the human brain by processing data through interactconnected layers, with each layer, refining information as it progresses through. I'll have a nice concrete example of that. These are the two that you see a lot that's being thrown out. They are hearing aids using machine learning, or we have some deep learning, that kind of thing.

To make this a little bit more real and concrete, we're going to use a fun analogy to differentiate between machine learning and deep learning. So with machine learning, imagine that you have an audiologists who follows the user realm, helping them to adjust their hearing aid settings throughout the day. And then over time, the idea is the audiologists can get better at predicting the user's needs. So there's a couple different stages, then they're involved. We have what's in the first panel here, data gathering.

So at this stage, the audiologists needs the user to tell them what changes to make in the different environments. For example, the user might ask to reduce background noise in a noisy restaurant, in a quiet room. They might prefer less amplification. The important thing is that the audiologists learns from these examples. And then we get to stage two, where they're learning patterns.

So as the audiologists spends more time with the user, they begin to notice patterns. They remember that every time the user's in a noisy restaurant, they lower the background noise. Similarly, they notice that in quiet environments, the user, for example, prefers a softer sound. And at this point, at this stage, the audiologists doesn't need to ask the user what changes to make every time they start predicting what the user will want based on their past experiences. And we get to the third panel where they're making predictions.

After learning enough patterns, the audiologists can now make adjustments automatically. They don't need to wait for the user's input every time. For example, when they see the user interact restaurant, they'll automatically reduce the background noise thumb being asked. And the important thing is, the more examples that the audiologists experiences with the user, the better they get at making these automatic adjustments. And then the fourth thing that's at the bottom here is improvement over time.

So if every new experience, the audiologists can refine their understanding of what the user's preferences or their needs are, and they'll continue to learn and improve as they encounter different environments and situations. And the idea is that they'll kind of get closer and closer to what the user needs without manual interactvention. I've heard a lot about deep neural networks. I'm sure you just may be familiar with the term DNN deep neural network. This is a little bit more abstract.

A neural network is like many, many teams of audiologists working together to analyze a patient's complaint. Each layer in the network, which is going to be represented by a team of audiologists, a group of audiologists with similar expertise. And then each individual audiologists would be like a neuron in that network. And what they're going to do is they're going to pass the job, the information, from one team to the next, and further refine the information as they advance through the layers, each layer adding their own individual expertise pointeract here represented. This is a classic neural network schematic where you have an input layer.

This is where the information gets fed into the network. You have an output layer where you have your options, there, possibilities that need to be selected in terms of what are you going to do with that information. And then you have a hidden layer. These are layers you don't see. And as indicated here, these layers can be in the millions.

Or if you're like Google, you'll have billions of these layers. So what's happening with these layers, though, is pretty interactesting. So each individual neuron, each individual audiologists, is passing different amounts of information to every other audiologists in the team. So that when you, by the time you get to the next layer, the individual audiologists, that individual neuron, is receiving a unique combination of information from every other person on the team. So while you have a team here, they're getting the same information, they're getting information from the previous team, but each individual person is getting different amounts of information, they're getting a different look at that information, they do their thing and pass it on, so on, so forth.

So let's kind of go and work through our analogies. The first team listens to the patient's complaints, gathers general information about their hearing issues, environment, the sound levels that they're in. The second team will focus on more specific details, understanding how the issues changes in different acoustic environments. For example, a third team, and they'll apply their expertise to filter out irrelevant information like background noise, and refine the data. And the final team's gonna combine all of that information and make a final diagnosis and recommend a particular hearing adjustment, hearing aid adjustment for that patient.

So it's just kinda like a neural network, each layer is adding information, refining as they go along. Each individual audiologists in that layer is like a neuron. And again, the important thing is they're giving different amounts of different kinds of information to every other team member. So as you progress, each team member is getting a unique look at that data. The important thing, too, though, is that it doesn't end there.

Once a decision is made, the important thing is it needs to learn to. Learning is the most critical component of this network. And this learning in a neural network is achieved through this idea called back propagation.

So where the information learns from its mistakes, or the network learns from its mistake and then tries to improve on it. So the way that this would work is that, so as we pass the information through, again, a decision is made. And if that decision, let's say, is not the right one for that patient, then what happens is they need to reevaluate. They need to reassess. What they do is they now work backward through the layers and the teams of audiologists and each audiologists.

Each individual is now looking at the amount of information and the type of information that they gave to every other team member going forward. Okay. And they're going to revise then, based on that feedback, they got it wrong. What do they need to provide each individual team member on the next layer? What, how much information do they give them to improve their decision making?

So, for example, maybe they'll realize too much emphasis was placed on background noise and not enough on the speech.

They'll do this iteratively. They'll go back through, make a decision. If it's wrong again, they'll work backwards. And what's happening in this learning process is the type and the amount of information that's passed along by each individual audiologists to all the team members gets refined. That's where the learning comes from.

Then the idea is, over time, the whole network, the team, becomes much more effective, much more accurate, and gets it right more often than not. And it does this iteratively thousands or millions and millions of times. That's how we train the neural networks.

I'm not going to go through the information in this table, but I'll just look at the key points here. The differentiator between machine learning, deep learning. Machine learning requires more structured data, requires more human interactvention. They're best for less complex tasks. The advantage with machine learning is they're more computationally efficient.

They can be done on small devices with not as much resources in terms of processing with regard to battery consistumption or power consistumption. Deep learning is more automatic, can deal with unstructured or structured data to deal with much larger data sets. It's better for more complex tasks. And of course, the downside is it requires much more processing power, more time and more actual power to generate, to power the deep learning. So the key takeaway for us is that machine learning efficiently focuses on solving specific problems of structured data, required human input to define the features.

It requires a little bit of human interactvention to figure out what's important in that process. Deep learning uses complex neural networks to process unstructured data, solve more complex problems with minimal human interactvention. And for us, in the context of hearing aids, that deep learning enables more sophisticated sound processing, while machine learning can handle simpler tasks like adjusting the settings based on the.

So where does all of this happen? Well, it depends upon which stage that we, that we want to consistider. And we're going to kind of take a more, a deeper look at this. So obviously, we know that they're probably done on super, this could be done on supercomputers. And to the extent that this is done in the hearing aid industry, this is in the background.

This is during the research and development phase where they're designing the algorithms, they're analyzing data that they've collected from users in the past. The cloud. So this be the cloud also is supercomputers. It doesn't live in the air, obviously, the cloud, again, it's just the only differentiator there is, how the information gets fed into the system. So this is via a, probably a wireless network, but again, still dealing with supercomputers.

The idea here though is it can get data from lots of different users. And the ideas that can help refine the hearing aid algorithms by learning from the individual users, but also aggregating information from lots of different users to improve the algorithms for each individual user.

While most cloud based AI improvements are made after the fact, not that it's not in real time during hearing aid use, they can provide ongoing algorithm improvements that can be perhaps downloaded later into the hearing aids for a more permanent solution for that user. Smartphones. The actual AI could be, to a certain extent performed on the smartphone, has much more processing power, much more battery power than a hearing aid, for example. They often will act as interactmediaries, so they'll process the AI commands and relay data back and forth between the hearing aid and perhaps the cloud or the phone itself. They can help manage non real time processes like sync the settings and updates, that kind of thing.

Then within the hearing aids themselves, this is actually interactesting, how is AI actually implemented on the hearing aid? The easiest of which is that the solutions that were derived using AI at another place, at another time during algorithm development, for example, noise reduction, to a certain extent, can be hard coded on the chip. So the AI was done at another time, and what's being put on the chip is the solution. All of those weights that we had talked about, all of the decisions about how the audiologists are going to consistider the information they're passing on, all of that's already occurred to, and they put that solution then on the chip itself. Or we could actually have AI performed in the hearing aid.

Of course, the limits are going to be is the amount of processing power that the hearing aid can provide on the chip with the amount of battery current that we want to devote to it, easiest of which would be machine learning, because, remember, they're computationally efficient, they don't take as much power then what I've called here, relatively shallow deep neural networks. I know it's an oxymoron. Did a little bit of research ahead of time. Shallow neural networks are actually that example I showed you where they're maybe two or three layers deep. That's technically a shallow neural network.

Deep neural networks are hundreds, thousands, millions, billions. What's on the devices to the extent that they are going to be not what we having the supercomputers, not what chat, GPT, or Google are doing, but a much more reduced, many fewer layers of orders of magnitude on the chip. So here might be for real time tasks like noise reduction and adapt the hearing aid settings.

Speaking of chat, GPT, everybody's heard. So I had just this fun acronym to kind of throw around, kind of summarize. Basically, you know, how is, you know, what's AI in hearing aids? So I called hat GPT. This is just a fun acronym, hearing aid technology guided by predictive techniques.

The idea here is this encapsulates AI, machine learning, deep learning, all of that, to the extent that it predicts and adapt to hearing aid users needs in different environments. Important thing here is using data driven predictions to optimize the sound processing, enhance the overall listening experience. And then the takeaway here, AI and hearing aids, exemplify by this broad concept, you love it or hate it, that's okay. It's transforming the way, though, that we as audiologists can deliver customized care by leveraging advanced data analysis and predictive techniques. And I venture to guess that's probably why most of you are here today.

How is this going to impact what I do? What can this offer my patients, what problems we have to think about what problems actually need solving that AI can provide a solution for. Let's take a real broad overview before we take a deeper look. AI and hearing aid research is not new by any means. We're almost a quarter of a century into AI being formally used in audiology.

And back in 1999 now came up with their first nonlinear algorithm, the Nana L one. And maybe to the extent you know this or you don't, but it actually used a primitive form of AI. They used an optimization procedure to derive their prescription rule. And basically what they did was they looked at a measure of speech intelligibility using the speech intelligibility index, which looks at the amount of audible speech information across different frequencies. So basically the idea there is how can we turn up the gain at different frequency bands to increase the AI, while consistidering the non linearities associated with the cochlea and with loud presentation levels, so on, so forth.

To a certain extent, louder isn't better, we know that. But also there was a control process pulling the gain down to consistider what was the perceived loudness of the amplified speech, so that it wasn't any louder than would be perceived by normal hearing individual at that same input level. So we got two things driving the gain up, but then keeping a cap on it. In terms of perceived loudness. They did this for, I don't know, like 55 audiograms, maybe more different input levels, okay.

To figure out what was needed to maximize speech intelligibility while keeping loudness no louder than for normal hearing individual. And push that through, figured out the solution, and then use basically another algorithm to summarize all of that information for all those audiograms, all those input levels into one prescription that we know of as the nail. As we forward then into the early two thousands, we start to get some rule based systems in terms of our classifiers and what we mean by rule based systems are mean like if this and this and if this but not this, then this. If the input level is high, the signal to noise ratio is low, then go ahead and implement aggressive noise reduction, that kind of thing. So things that we as people are kind of putting into the system if based on a bunch of contingencies and using that also in terms of, and basically in the context of our classifiers that we know and love, that we can talk about in a little bit here.

But then as the knowledge base and the processing power got to be greater, actually using data now, data informed models, that is database machine learning systems, where you're pushing things through at a much, with much higher samples, much more inputs. For example, we have some self tuning, then we get to the later 2000 teens, if you will, where we're using, integrating smartphones into this for things like health monitoring, engagement with the user. We have some R and D. We'll talk about quality control. And then where we're at now, cloud based AI, the deep neural network stuff, okay, being used.

The applications we're seeing being used most is for speech enhancement, noise reduction.

So we're going to kind of take a circle around the idea, the different ways that AI might be implemented in hearing aids. And this is ideal because I don't note to the extent that this exists in a single hearing aid, probably different amounts of this to different extents, how we can implement AI throughout this continuous process of improvement. And we're going to start with R and D research and development. So the easiest case here is what's happening behind the scenes. Things that maybe that we don't see that maybe the user where they're the beneficiaries of the AI, but the AI has not, they're not being, interactacting with the AI.

So this is the idea of developing the algorithms. So now researchers can train the algorithms that are the features in the hearing aids to understand complex listening environment by feeding it lots and lots of samples to tens of thousands, to the tens of millions, okay? And then based on that, allowing those networks to figure out the solutions, if you will adapt, create adaptive algorithms that can automatically adjust to hearing aid settings based on the acoustic environment. I'll give a more concrete example. This is rather abstract, I can appreciate that.

Again, deep learning models can be trained on lots of data sets to improve real time algorithms. So the idea here is we can create the algorithms using this data, but then also create algorithms that are continually improving in real time use, that are separating, say, for example, speech from noise. The other interactesting way is that maybe, again, things we don't see on our end as end users. But using, I saw one manufacturer talking about using the Google speech, or yes, speech to text engine as a metric, and they basically looking at how accurate is it at predicting human performance in terms of speech understanding with a particular setup. I believe they were investigating t coils in this use case.

But the idea there is, can we get a better way of estimating real human performance using some AI behind the scenes? Lots of people have been for a long time and from our research world have been looking at how we can predict human performance using different models that mimic the cochlea to varying extents. But the idea here is to refine our algorithms without actually having to bring in, you know, hundreds of people to test them on. So before we actually do get to that phase of testing, we've already have a good idea about what's going to work and what's not going to work by using some, again, AI machinery to predict human performance for these different solutions. Okay, so to make this a little bit more concrete, this is an example from one of the manufacturers.

They offer deep neural network noise reduction. And what they had done was they recorded many, many, many real listening environments, noisy environments, quiet environments, so on and so forth.

What they did was they used a high level of sophistication by using a microphone sphere, so consististing of equal arrangement of 32 microphones in a 3d space. Here we have an example on a subway plex. And what they did with all those recordings, they went back then into the sound booth and made recordings of, they put a circular array of loudspeakers.

So the idea here is, can we recreate this acoustic scene? Well, first you need to know how each of these microphones is going to react to each of these individual loudspeakers. So you could basically invert the process and create this realistic 3d environment.

And then what you can do then is do the same thing, but using, based on where the hearing aid microphones. So putting a mannequin in that same acoustic space and figuring out how that sound coming from each of these loudspeakers is going to arrive at the hearing aid microphones. Okay, for different styles of hearing aids, whether it be above the ear in the year, that kind of thing. All right? And so then what they did was then they replayed, played back through using open headphones, the recorded scenes that they have basically invert the process from here so that a person was actually listening to those noisy scenes.

But the key thing here is overhead phones, because what you see with some of these algorithms that they'll do is kind of in the research world, they'll go to YouTube, they'll find a bunch of examples of noisy backgrounds, and then they'll go to a classic speech database that's been recorded in quiet, and then they'll mix those together and have their speech in noise. Well, we know that when we go into real noisy environments, we don't speak the same way as we do in a quiet environment like this. We speak differently. We have our vocal, as our vocal emphasis goes up. Okay, for example, we put more energy towards the high frequencies, that kind of thing.

So the idea here is to recreate how this person would be speaking in that noisy environment because they're listening over head, open headphones to those noise scenes. So they're going to adjust their speaking, their vocal effort accordingly. And then we record them. And the key thing is that they're over headphones. So what's being picked up at the microphone here is just a clean speech from this person, but they're speaking as if they're in the noisy environment.

But we're not recording the noisy environment, just their clean speech. And the key thing then is then later in time, we can take those clean speech samples, as we indicate here, and mix them in with the speech and noise. So it's basically creating a realistic environment where there's speech and noise. And then we have the clean speech. And then what they did is fed that into a neural network.

So we have a bunch. So you have the same speech. In this case, it's clean. In this case, it's mixed in with the noise. And what they try to do with their neural, deep neural network is figuring out how they needed to adjust the gain in individual frequencies and time instances so that the speech and noise looked as close as it could to the speech client.

And that was the solution then the deep neural network noise reduction solution that you could put then on the chip. So they figured out what the solution was, put the solution on the chip. So the chip maybe may not be actually doing AI per se, but utilizing the results of the AI to help then improve speech and noise.

One of the things, too, that's actually interactesting in this process is that what I've learned is that the system that they've developed actually utilizes directionality and beamforming to help get a better estimate of the signal to noise ratio, individual frequencies and times. So they can then differentially apply the gains while the user is actually in a different, maybe in a different direction of mode. For example, the patient can be in an omnidirectionalal mode. At the same time, the hearing aid can be using the information from beamforming, from directionality to help inform the algorithm. So it's happening behind the scenes, even though that we, while we can present the patient with a different directional pattern.

And the idea here is this goal is to achieve equal improvement in terms of signal to noise ratio and speech intelligibility. That would be equivalent to directionality while presenting that patient perhaps with an omni directional response or maybe something that's not as beamforming, for example, just a classic directional pattern. Because the thing there is, is the idea is that beamforming has its time and its place. You have a very narrow focus of directionality. But the problem, inherent problem, is that if the thing that the person wants to hear is outside the beam, they're going to be cut off from it.

Or in a dynamic situation where the conversation shifts, could the dominant conversational partner can shift very rapidly. The individual who's beamforming is going to miss out on the change of conversation, and they're going to be searching around and people will be laughing at the table, and they're going to want to know who or what. What's the gossip about? Because they've missed the chain, the handoff in the conversation. So the idea here is you could present that patient with a broader directional pattern and allow the deep neural network noise reduction to handle the rest without having to put them into a narrow focus of directional.

So this is an example, one example of how we can use AI in algorithm development in the research and development phase. And then those solutions, in an ideal world, perhaps, can get fed into the automatics. And we're all familiar with automatics. A lot of us have cut our teeth with the automatic classifiers in hearing aids. This is not new news, but let's take a few steps.

Think why we need classifiers. Well, we know that hearing aid users frequently move between different acoustic settings, quiet rooms, noisy streets, etcetera. And if we didn't have classifiers, we, the patient, would need to make manual adjustments. That's going to take some cognitive energy to figure out, okay, what do I need to do? Is this a time to switch?

How do I switch? They have to be able to physically make the switch. Okay, not going to be easy to use. And if you're having to do this frequently, patients eventually going to fatigue and just stay in one mode.

Let's think about what classifiers do. They're probabilistic, so what they're doing is they're making a prediction. They're sampling the environment and trying to generate, basically, what's the probability based on the acoustics that the patient is in any one of these environments? Because it doesn't know for sure. It's not 100% sure.

It just has more likely than not. And that's the solution that's often chosen. Which one is the most likely sound scene for this patient based on the information that it's taking in. And then it makes the appropriate switch as pre designed by the hearing aid engineers. What to do when you're in that sound environment?

Of course, as we know, as audiologists, we can get in and customize those often these individual scenarios, just the amount of directionality, the amount of noise reduction, for example. But then it's a hard coded switch over again. Refer back to the timeline over the years, though, that these classifiers have been trained using, to different extents, AI. Can we get a better predictor of what this sound scene is other than what we think as humans are the deciding variables? Because, again, we can only take in so much information, information at so much time.

And we try our best to, if this and this and this and not this, then this a rule based system. But we can greatly expand the amount of variables under consistideration if we allow AI to do this to a certain extent. Not only that, not only are the classifiers themselves trained up using AI, again, by presenting it with lots and lots of examples that maybe we've pre classified. Okay, here, AI, this is all speech and noise. This is speech and quiet.

This is some speech, maybe with a little bit of music. Okay, you get different combinations. You basically tell the, tell it, feed it a bunch of examples, tell it what they are, figure out, have it, figure out what it needs to do to reliably classify those. Not only that, then the response of the AI itself can be, um, trained using AI. Okay.

So not once you make a decision what we're at, how do we respond? Could also benefit from some of that advanced computing. All right, so kind of went over this, again, early, classifiers being rule based, limited by our knowledge, and again, machine learning algorithms, neural network algorithms being used to refine our classifiers. An important thing here, though, is that the advantage we got to have to ask ourselves, is there a problem that needs solving? And I think we can all argue yes, because what we're going to talk about in a moment here is that classifiers can and do fail.

But the thing here is, if we can get a greater data set, what we want to do, be able to do is have these classifiers be more accurate at predicting less common environment, maybe environments that they were never trained on to be able to adapt to those as well, because they've had that vast training experience to be able to get to the more narrow individual situations and solutions that we couldn't otherwise do using the old system. Here we go. As I indicated, I emphasize acoustic. So, just simply based on the sounds coming in, the hearing aid classifiers using acoustic information can and do fail. It can fail to accurately capture the listening environment.

And there's been at least a decade, decade and a half of studies. The easiest ones are taken, a hearing aid from each of the manufacturers put it in an identical listening environment, and, you know, in a controlled acoustic situation that you know what it is, and you look at the data log and see how the hearing aids have classified those environments. And you look and you see that these hearing aids have classified these environments differently. The speech and quiet ones are easy, but once you get the more complex environments there, you start seeing differences across the manufacturers. So even though that we have all these advanced solutions, each of the manufacturers, maybe to a different extent, are coming up with different probabilities of what that actual environment is.

Okay? So to a certain extent, they're failing to actually capture the listening variant. What we'll talk about is, even if they get that right, even if they exactly know exactly what environment they're in, the user might have different listening needs, and they might have different listening intents. So let's take a look at that. So an example of different listening.

So we look at this busy restaurant situation, and we might argue that regardless of where you're at, you're going to get similar acoustics. Let's just pretend that we have two individuals. Same hearing loss, same hearing aid, same acoustics. Hearing aid gets 100% accurate classification. This is a speech and noisy environment for this individual here.

If he's wearing hearing aids, he's going to want, he's trying to carry on a conversation with his dinner, so he's going to want more aggressive noise reduction to turn off the conversations, reduce conversations going on behind him, and a more directional focus. The waiter, on their other hand, same hearing aid, same hearing loss, same acoustics. He's going to have different listening needs. He's going to benefit from a more broader frequency, an omnidirectionalal mode, to be honest, because he's going to want to be able to hear everything going around him as he navigates through the restaurant, okay? He's also going to, you know, to be able to have some vigilance of things going on around him, people walking behind him so it doesn't bump into him, but also people calling for his attention because they need something.

He needs to be able to have surround sound, okay? Same acoustic, same hearing loss, same environmental classification, different listening needs. And so this kind of demonstrates this balance in terms of what we need for effective directionality. So, effective directionality requires careful balance between directional focus on the one hand and environmental awareness on the other hand. So if we revisit our restaurant vignette, recall that our patron will benefit from some degree of directionality when communicating with this dinner companion across the table.

This might be in the form of traditional directionality, which we'll call monaural beam forming here, or binaural beam forming, where we can get the much more narrow focus. But also recall that the waiter requires a certain degree of environmental awareness in order to create a spatial mental spatial map. This can help promote the sense of psychological comfort via the increased environmental vigilance that's afforded by what we can call surround sound, and also this ability then to hear for people calling from for his attention from different directions. And the ability then to create this mental spatial map is going to depend upon preserving or recreating pen acoustics, which provide the monorail cues to help us tell sounds coming from the front to the back. In addition, it will require that the user have access to low frequency interactactions, time differences, or itds, and high frequency interactval level differences ilds, which then together then provide the auditory system with the binaural cues that it needs to help create this mental spatial.

Both of these, then are also mediated by the acoustic coupling to the ear, which refers to the relative position of the microphones above the ear in the ear, to which extent the ear is open or occluded, and whether not, the signal processing between the two hearing aids is integrated. For example, independent signal processing in the left and the right hearing aids might distort the binaural cues, especially the intro level differences. All of these need to be consistidered for effective functionality to meet the patient's listening needs. Let's take a look at this busy Nashville bar scene. So what we want to emphasize here is different listening intents.

In this environment, too, individuals can be exposed to the same acoustics, same hearing loss, same hearing aids in the same environment, but require different signal processing solutions if they have different listening intents. Take a look at these two individuals who are trying to carry on a conversation.

They want to reduce the conversations going on around them. That music that's playing on the stage, that's background to them. They want to be able to carry on a conversation.

This individual, he doesn't want to hear their conversation, nor anybody else's. He wants to focus on the music at the front, different listening tense, everything else being the same. So even if the classifier gets 100% correct, there's still room for improvement. So one way we can do this is by giving these classifiers more and more information. How does the information get acquired?

Well, through sensors. The easiest sensors are our acoustic sensors, that is, our microphones, our dual microphones on the hearing aid, and then to the extent that we might utilize binaural processing across the two hearing aids to get additional information. We have acoustic sensors that's very powerful indeed, but we can certainly, there's getting room for improvement. Motion sensors. These include accelerometers, measure linear acceleration, detect changes in speed or direction.

Think of our vestibular system. Gyroscopes. This can rotational motion or angular velocity, and then magnetometers. Magnetometers, these can use the directional of the magnetic fields in the earth, much like a compass. This is kind of in a lighter shape because I don't know too much that this is implemented in hearing aids.

Okay, but accelerometers for sure. I know at least when manufacturers talk about using gyroscopes, the point here is that then where's the rule for AI? We include extra information. This is, again, more information we add, okay, we're going to start running into limits in terms of our human ability to be able to know what to do with that information, to be able to analyze that information. So we can use AI to help develop algorithms that combine the data for multiple sensors from the microphones combined with emotion sensors to better understand how that user is engaging in their environment.

And we also can then also get a better sense in terms of the reliability and the predictability of this information. Is there any value added benefit? Well, we can provide lots and lots of examples, big data sets, just understand patterns of behavior and different environmental conditions to see if there's actually something there to be had. Is there something reliable to even gather from these different sensors? And again, the idea there is this is vast amounts of information that it would be difficult to comprehend.

The solutions that are offered are combinations of different aspects of the input acoustic motion that maybe we, as humans cannot visualize. That's usually often what's, that's almost always true of these neural networks. So you peel it open, you're not going to be able to see the answer.

Then the hearing aids can use the sensor input and then the AI in the hearing aid to know what to do once it figures out what's the condition, what's the environment, how do I respond? How do I respond over time? For example, by monitoring the acceleration forces along three or five planes, x, y, and z heredity, an accelerometer can provide the hearing aid with a rich set of information concerning the user's movement or lack thereof. For example, this figure in the bottom here shows motion patterns for three scenarios. So, person sitting at rest will show small amounts of amplitudes in the total acceleration or g over time.

In contrast, a person running will show large amplitudes and periodicity over time, whereas a person riding a bike may show smaller amplitudes with less periodicity over time. And all of this information then can be used to indicate how the user is engaging in their environment or that their auditory needs have changed. So we go back to our restaurant scenario. Accelerometers could detect that this person is sitting, they're carrying on their speech, speech to the front, poor SNRs. This person is going to benefit from some sort of beamforming detect that this person is moving around in this environment.

Open it up the scenario. Open up the surround. This individual gets up to pay for his check his senses. There's now been a change in his. In his motion.

Now these in motion automatically switch from the narrow directionality to the omnidirectionalality. So to this center, individuals walking when we walk. When we're walking and carrying on a conversation, we're usually not walking some other persons in front of. Usually we walk side by side with our companion, and the conversation is coming from the side. So here we're going to benefit from a focus from one side more to the other, or just simply an omnidirectionalal mode to be able to pick up on that communication.

You're hearing a lot about intent, listening intent, and being able to incorporate greater amounts of information using head movement. And this is going to be very vague, very broad. This is, again, kind of limited to what we can conceptualize through our examples here. But if we just think of some very basic examples of how head movement, not just simply linear acceleration, can indicate how a person is engaging with their environment. So if a person's turning their head towards a sound source, such as a dinner companion and a noisier restaurant, obviously that might indicate a desire to focus on that conversation over the other sounds.

So the solution may be to implement a more narrow beam of directionality is depicted here. Or maybe it's not speech, maybe it's something else. In this case, this gentleman here is trying to listen to the jazz pen. So trying to infer from just a simple turn of the head, how about a busy table, family gathering, thanksgiving, everybody's getting together from different places, haven't seen each other in a long time. Conversations are flying all over the table.

Dominant conversational partner shifts very rapidly, right? So you're consisttantly maybe turning your head from one to the other. Okay, you're trying to search, maybe scan the environment, figure out who or what the thing is you want to listen to. Maybe you lock in. I want to hear this person not that person.

Okay, so the idea is, you know, if you're turning your head a lot, maybe that would indicate you need to broaden the directional focus so that the person can actually analyze their environment and they stop. Maybe implement temporarily a more narrow beam of, or just simple directionality.

How about leaning your head in or tilting? If you lean in a little bit and it's noisy, that could indicate that the person's really working hard to understand the person in front of them. So maybe that would signal to the hearing aid. You need to implement a more aggressive noise reduction or directional solution. All right, so those are the automatics.

Let's get on to what happens. How the hearing aid can might learn from how the user's interactacting with their hearing aid. We have stuff that could be occurring behind the scenes that requires no user interact. Well, it requires interactvention, but the AI is not directly interactacting with the patient. They're learning from how the user is adjusting their hearing aid over time so that they can better fine tune the settings for when they're in similar environments.

The idea then is as the hearing aid learns from the user's adjustments, so it's not explicit, okay, they're just kind of sitting behind the scenes watching, observing, trying to refine its own algorithm over time. Then we have those situations where, okay, the automatics have failed, they've done their best, maybe they've accurately classified it, but maybe the solution, okay, maybe the solution isn't the right thing for that patient in that environment. So here in this case, the user could work with AI to find an optimal solution. So, through an app or just a hearing aid interactface, the one I have in mind is one where you get the user uses their app, the phone on their app, to engage the AI. And AI provides them with a b comparisons, two different programs, a b, the patient listens to both, and then they rate how much they prefer one over the other.

That information gets fed through and it presents them with another comparison and does this iteratively over those few minutes until it figures out what that patient needs in that environment. So again, maybe the classifier got it right, but the solution implemented is wrong. So the AI learns get better at predicting it. So next time they're in that environment, it doesn't need to go through that process or doesn't need to go through it as extensively. And then they also, in this particular use case, the ability to collect information from other users can go up through the cloud so that then each individual user can benefit from others who have been in those scenarios.

And then we have instances where the user engages the AI, but there's no cross interactaction, it doesn't require any further action by the patient. So then in this case, the hearing aid is going to implement a much more advanced solution. It might happen to a certain extent on the phone or to the hearing aid itself, where it just kind of does a much deeper analysis of the environment and then implements a solution. So this could be a tap on the hearing aid that indicates, hey, I'm in trouble, do your thing, or opening up an app and doing the same thing. And then that brings us to analysis.

Here we are, analysis. So again, analyzing data sets on an ideal world, that analysis, the data from the user interactventions, could get fed again to the cloud, which eventually maybe get downloaded onto the company computers. And then the idea is now you got all kinds of data. How do you analyze this data? How do you make sense of it?

AI could be a perfect way to make sense of this vast amounts of data to help refine the next iteration of the algorithm, or maybe something that could get implemented as a firmware update, or maybe the next generation of algorithms. One of the manufacturers talking about using AI for quality control. So by monitoring streaming stability, battery performance, hardware integrity, trying to infer whether or not the performance of the hearing aids going down over time, also, you could use this information maybe to detect more systemic problems. So if a bunch of users are experiencing similar degradation in performance, this could indicate that there is a problem at a larger scale. Overall.

The idea here is to use AI to ensure hearing aid longevity and reliability. Some of this might be, I'll go through this relatively quickly. AI and health monitoring, we know with different use cases of being able to integrate with the smartphone, keeping track of the patient's physical activity, their steps, their social engagement, how many conversations they're having, fall detection. The idea here is maybe the sensors, the accelerometers on the hearing aids are more sensitive and so specific specificity and sensitivity to falls, not other things. Tracking their activities, using the app on the hearing aid, on the smartphone, as a hearing health coach, much like some of us might have for our healthcare providers, that encourage you to eat better and to exercise more.

Same concept applied to hearing aids. It's not for everybody, obviously, but the idea here is to encourage them to engage in more social conversations and to maybe be more physically active. That kind of thing also provide information for the audiologists to look at.

And the whole idea here is that to the extent that our hearing aids connect with our phone, our phones, our smart systems, it kind of gives this added aura to our hearing aids being smart, even if we are just simply connecting to our hearing aids. So we're using our, our voice assistants, your Google, your Alexa, so on, to do the same things we would do otherwise. connection to the cloud again offers this extra aura of smart sophistication. So things like real time language translation kind of get a summary here of just the different ways that have talked about how AI is proving hearing aid functionality. Here's a summary of some of the different features we've talked about.

And I told you at the onset that's not likely you're going to see all of this in one hearing aid, but you'll probably see some of this to varying extents across the different manufacturers. And to bring this to a close, what's your role in all of this? So you have now have tools to help refine and optimize hearing aids based on ongoing user feedback and real time environmental data. Important take home message is that AI can support, not replace, your expertise. They can automate some processes, but your role is still critical in the process.

You as audiologists, as the trained audiologists, are essential in guiding the use of AI for individual patients, interactpreting those results, providing some context, and ensuring the best outcomes based on their unique individual needs and use AI as a collaborative partner, can compliment again, not replace your clinical knowledge again. AI can manage ongoing adjustments and collect data. But the human touch, your human touch is key to understanding the patient's emotional, social, psychological needs. That's why we go into this field. This is the thing that we take home at the end of the day, and AI cannot replace that.

And with that, I'm going to wrap it up and I will take, see if we have any questions. I see there's still good attendance. So I do appreciate that there's no questions at this time. I be happy to take any or any right now, you know, and if anything, you know, kind of as I tell my students, I teach the hearing aids courses here at Purdue. I get the students for a year and one of the things I tell them is you're really the most important thing you could do is stay informed.

This is a rapidly changing area. It's breathtaking. And to be aware is we all have our preferred brands that we're going to dispense, but be aware of what's going on around you, even with the brands that you don't dispense, so that you can be better informed when your patients come to you with some of these specific questions, what to expect when others question what to expect will happen with the launch of Apple and AirPodss Pro two s. I think the consist. Well, I like to think, I like to think the consistensus is that maybe not a consistensus, but I don't think it's going to change a whole lot, to be honest.

I mean, we still have all of the individualization. I mean, that's just this theme of this whole OTC business, right. I guess the only aspect is that Apple brings the name recognition, maybe some assurance of quality, maybe some assurance that, hey, these things don't work as hearing aids. Fine. I still have great earbuds.

I really don't think that they're going to be. People are going to where. Where those around as full time hearing aids. I could maybe write battery power has to be a limitation. Those things can't run all day like a hearing aid can, but maybe in different select uses, temporary cases.

But I think somebody's pulling out their airpods for different environments. They're going to eventually start to feel that they need a more long term solution that maybe hearing aids can provide. Next question. We've seen AI play a role in physical impairments such as poor balance. Are there any developments in AI involving assisting in cognitive impairments like temporal processing?

I'm sure there is, but to be honest, I don't know. I can't give a specific answer. But I'd be more than likely to bet that somebody is working on that behind the scenes, particularly. I think the value there I think we can appreciate with AI is the individualization. So you just think like my kids, you know, young school kids, maybe some college age kids, you know, the ability of AI to adapt the tutoring, the lesson plan for the individual wherever their levels at.

And I think what if you're thinking about rehabilitation, oral rehabilitation, cognitive rehabilitation, it's the idea that you could have a very patient assistant on the other end, and that could tailor the lesson according to your needs.

Any comments regarding AI in patients with visual impairment? That's probably just. I'm just. I found that they perform more omnidirectionalal. That would be interactesting, actually.

That's. That's interactesting.

Yeah, that would be true. I think I've seen that literature where people with. Yes, I know what you're referring to. We came out recently, they reduced vision, that they're going to prefer to be more in an omnidirectional. I would wonder to what extent you just keep those individuals in omni all the time, but maybe with an occasional override depending upon their needs, because they're definitely going to need to navigate their environment.

Maybe a new situation that I think about it would be, is using those accelerometers, right? So if they're sitting and they're not in motion, that would probably be a good instance where you might be able to have the more directive directional microphone response. But obviously, for that person, anytime they're in movement, they're going to need that full surround. My favorite use of AI in hearing aids. Oh, that's fun.

I'm really my favorite. I guess my favorite. I don't know, maybe my favorite, but the one I'm most excited about is the ability of noise reduction to provide actually a signal to noise ratio improvement that before could only be achieved with directional microphones. Because for the longest time, we know that everybody has been looking for noise reduction to have an improvement in signal to noise ratio and never being able to do so. And this idea of kind of giving an extra value added benefit, you could provide some of the same benefits as directionality.

And remember the teeter totter where we had this idea? You need environmental awareness and then the benefits of directionality, that's kind of a way of evening that out a little bit without having that trade off. So I'd say that would be mine. I'd like to see hearing aids eventually use voice recognition like Alexa. Identify which, oh, which voice you want to hear in background noise.

This thought has been conceptualized where you could perhaps use a cloud based solution to train it to identify preferred conversational partner, spouse, or somebody you communicate with frequently. So then when you go out in those environments, you know, you get to that mode, you know, wife mode or something, where then it's going to preferentially pull that voice out from the other. So I've heard of that conceptually. If it will ever be ready for prime time in hearing aids, who knows?

Oh, any one manufacturer that has more advanced AI. I tend to shy away from giving specific recommendations for any reason. As my role as an educator, I like to provide people with the information and have them make an informed choice, and it'll depend upon your patient, individual patient needs. I think all of us can agree that the ones that we choose, at the end of the day, sometimes it's the other variables, whether it be customer support, the ease of use.

My easy answer to that is it's probably just a Google search away. You're going to find two that pop out above the others. I'll leave that at that. Well, I don't see any more questions. I'm thrilled to see that we still have a high participation rate at this point.

I will go ahead and kind of bring this to a close at this, but feel free to engage in any other way following this talk here and bring up the last slide. There we go. And I'll hand it off from here. Thank you so much, Doctor Alexander. Just a reminder, everyone, for those of you who wish to earn CE credit for this course, please just complete that multiple choice examined.

Once this course has ended, we like to thank Dr. Joshua Alexander for his time and expertise. We hope to see you all back soon. Thank you.